R/D 4624-CC-01

Contract DAJA45-85-C-0014

FINAL REPORT

September 1st 1987

"TOOLS FOR LEARNING"

Yves Kodratoff
Equipe Inférence et Apprentissage
UA 410 du CNRS, LRI
Université Paris-Sud, Bât. 490
91404 ORSAY cedex FRANCE

The main research topic of this contract, i.e. testing how symbolic Machine Learning (ML) techniques can be applied to Scene Analysis, has been completed by the final implementation of our system that learns features to recognize multi-funt characters.

The work done during these last months has been a preparation to the application of these techniques to "real world" problems, in which raw noisy images are provided to the system.

The approach we want to develop can be summarized as follows:

1- Inductive Learning of decision trees from a fixed vocabulary.

Such techniques exist (Quinlan *Machine learning Journal 1*, 1, 1986) and are currently under refinement in the ML community (see for instance, "Progress in Machine learning, I.Bratko&NLavrac eds. Sigma Press, Wilmslow 1987)).

2- Learning new descriptors improving the decision tree.

We have developed such techniques, they have been reported in my last report.

3- Merging the results of 1- and 2-. Discriminant methods.

We achieved a program that develop a new decision tree each time a new descriptor is introduced. The resulting methodology can be qualified of "discriminant method" since it learns to discriminate the examples of the concept to recognize from instances of the other concepts.

We are not yet able to apply such discriminant methods to any kind of noisy data. The choice of the descriptors that minimize the noise has been studied. The copy of our IJCAI-87 paper on the topic is joined.

4 Finding intentional recognition functions.

Being given a set of examples, the problem *is to find a* recognition function that describes them "in intention", that is which contains all their significant features. One must beware that such a recognition function does not need to be discriminant. One is in a situation where one knows to which concept belongs the instance to be recognized, but one wants explanations on why such instance belongs to the concept.

We have been building systems that create such recognition functions. As an instance of our results, find a paper that is accepted for publication in a forthcoming issue of *Artificial Intelligence Journal* (to appear in Summer 88).

Up to now, all efforts have been done toward the generation of discriminant recognition functions.

In our opinion, this explains the failure of all existing systems to be adaptative enough for coping with unexpected situations. The "discrimination" (recognizing that an instance belongs to such concept) and the "recognition" (explaining why it does) are distinct tasks of different level of conceptual difficulty.
For instance, recognizing that a given set of spots is a "human face" does not mean that a spot is recognized as an "eye", with the functionality of vision.

5 Merging "discrimination" and "recognition".

There is nothing done on the topic as far as our knowledge goes.
Our next proposal of a research contract to the US Army will explain how will shall work on this problem.
Simultaneously, we will adapt our existing tools to the problem of recognizing a human artefact in a landscape. During my visit to Fort Belvoir (to Dr Lynn E. Garn, June 18th 1987) I could see that our main problem, viz. finding the place (or window) to look at in the whole landscape, is already solved by his group. This makes a future collaboration bearing strong hopes of actual results.

For your information you find also joined a copy of my report on the Symposium on Pilot Aiding held at Irvine just after my visit to Fort Belvoir.

(Read) next letter

(35138) 88-02-05 12:09 Marco Valtorta CEC-DGXIII Comment to: (35028) by Yves Kodratoff CNRS Receiver: Marco Valtorta CEC-DGXIII Receiver: Yves Kodratoff CNRS Subject: travel permission ------------------------- I havee some problems with supporting a trip by two people to go to the same place. Please give me some justificcation for that, or accept funding for only one person between you and Rouveirol.

(35138) comment in (35140)

What do you want to do? (Read) next letter, (Join) next conference, Wait (for news), Comment (on entry), Personal (answer), (Re)type (the message), Review (entry), Quit, (Get) help, Other. -
(Read) next letter

(35140) 88-02-05 12:13 Marco Valtorta CEC-DGXIII Comment to: (35138) by Marco Valtorta CEC-DGXIII Receiver: Yves Kodratoff CNRS Receiver: Marco Valtorta CEC-DGXIII Subject: travel permission ------------------------- conditions: be sure that the consortium agrees, write a report.
(35140)

Report on the Symposium on
Pilot Aiding and Machine Learning (ML)

Held at University of California at Irvine (UCI), June 27th 1987
(Kodratoff's travel to the USA, 18-28 June 1987)

Yves Kodratoff
Equipe Inférence et Apprentissage
LRI, Bld 490, Université Paris-Sud, F-91405 Orsay

FOREWORDS

Since Manago and Ganascia are going to report on the International Machine Learning Workshop (which was our primary interest) , I will concentrate on the Symposium on ML and Pilot Aiding that followed, and that I could attend. Nevertheless, I will do 3 statements about the Irvine IMLW.
Firstly, some statistics that speak by themselves. Over 180 participants, 66% were US academics, 11% CEC academics, 4% Australian academics, 3% academics of other countries, 16% were US industrials which, of course, leaves 0% for any-other-country industrials.
Secondly, this workshop, from which stemmed the AI approach to ML (the book Machine Learning 1 is the proceeding of the 1st of such workshops), will stop being on invitation only, and start being an open congress, held each two years, starting in 1988. John Laird (Univ. Michigan) will be its first Programme Chairman. As a direct, even if modest, result of my action as leader of the COST-13 project on ML, it has been agreed that the congress will be international. There will be 3 European members of the Programme Committee, Programme Committee meeting will take place on the Eastern coast to ease European's travels, and the 1990 meeting will be held in Europe, being understood that the Australians will also take their turns in organizing the congress. I must add that most Americans are aware of the European worries about European under-representation, and helped me a lot in convincing everyone to start something that will be really international.
Thirdly, and of less world-wide importance, let me recall the delicious orange juice and yogurt at UCI restaurant and the general good quality of the food (there has been so many disputable complaints about the food there, that a little word of praise may not be totally out of place!).

ABOUT THE SYMPOSIUM

It was organized by Honeywell and the University of Michigan with the obvious help of several agencies. Honeywell has already started a thorough study of Expert Systems (ES) in Pilot Aiding. Due to the topic high variability, they have been naturally drawn to ML techniques. In order to get more information about ML abilities to cope with their problems, the existing team started this symposium were some of the US leaders in ML have been invited to brainstorm on the topic.

The Honeywell team is well-aware of the state-of-the-art in AI in general, in ML in particular, and some have been attending the workshop just before. This Symposium was an almost perfect prototype of the way industrials trusting their academics should start ambitious collaborative projects.

The aim of the Honeywell team is the automation of the crew station, with an emphasis on co-pilot simulation. They want to apply ML techniques to help them solving some of the typical difficulties of the problem: uncertainty, time, presence of external agents, changing environment.

PRESENTATION OF THE PROBLEMS

Valerie Shalin (Honeywell) presents the problems met when building an ES in pilot aiding.
The system must observe the situation inside and outside the aircraft, coordinate with the other aids to the pilot, take into account the possible errors in the information they provide, monitor the displays, and help the pilot taking the appropriate actions. Two major features make the problem more difficult: the system has to be real time and low tolerance errors.

Paul Scott (Univ. Michigan) delivers a talk on the classification of ML systems and techniques. I obtained a paper which still under revision. When I will get a non restricted version of Scott's paper, I will distribute it to COST-13 participants.

He describes the important role of taxonomies in ML, and presents his own (meta!) taxonomy of ML research. He makes a deep difference between a functional approach that concentrates on what changes improve performance and a structural approach that builds representations of experience. Learning involves two different kinds of searches (the topic: Learning as Search is a famous one in ML). One is a search within a representation space in order to find a good representation of knowledge, the other is a search through an experience space in order to find informative experiences.

Ed Wisniewski (Honeywell) insists on the importance of updating and refining knowledge bases seen as a ML problem. He proposes ways for the evaluation of the ML systems, e.g. their learning domain, representation schemes, possible experiences, experience generator, representation generator, representation evaluator, values of representation etc ...

John Laird (Univ. Michigan) presents the SOAR philosophy in ML. SOAR presents only one learning mechanism, and a rather weak learning component. Its learning power comes from the power and diversity of its problem solving unit.
He starts a discussion on the possibility to start from a given ES and "add it some learning". A discussion with Michalski and Mitchell follows. Michalski broadly approves of Laird's opinion, while Mitchell makes the point that he needed to write a translator between the existing ES in VLSI and its associated learning unit: LEAP.

Jim Anderson (I missed his affiliation) presents the connectionist approach.

Ryszard Michalski (Univ. of Illinois at Urbana) presents a discussion of Scott's framework for ML. His main 3 critics are that Scott does not properly take into account interactive learning, the cognitive constraints and learning by analogy. Scott has been classifying learning systems rather than learning strategies. Learning strategies include differences between rote learning (learning by copying from the strategy point of view), learning by being told (learning by restructuring knowledge), learning by deduction (i.e. by using specialization strategies), by analogy, and by induction.
He also adds some additional critics on Scott's description of Michalski & al.'s systems CLUSTER and AQ11.

Pat Langley (UCI) describes four strategies for learning. Chunking (in a broader sense than the one used in SOAR) means finding PART-OF relations, is an essential component of the discovery systems. Clustering, i.e. finding INSTANCE-OF relations generates extensional definitions. Characterization generates intensionnal definitions, often called also generalizations. Organization of the first three strategies allows information retrieval.

Valerie Shalin (Honeywell) presents a too brief description of the domain analysis of pilot aiding. She finds three main problems: classification, interpretation, and monitoring. Each applies to five main domain components: system status, situation assesment, tactics planner, mission planner, psychological problems.
She characterizes also a set of typical learning problems: classify non fused features, optimize emergency procedures, coordinate existing emergency procedures to handle similtneous emergencies, for example.

PANEL DISCUSSION

Michalski sees three problems. Merging new ML-acquired knowledge into classical ES already obtained. Analyzing the data, and finding regularities. Improving efficiency, through EBL for instance.
Mitchell thinks that ML is high risk/high stakes topic. Hard problems like "reading between the lines" of a repair manual are not likely to be solved in a near future. If domains have a well-known theory ML can be of immediate help, otherwise he doubts of it.

WHAT ML CAN BRING TO PILOT AIDING?

Valerie Shalin recalls the problems addressed. The future system should be able to fulfil four tasks. The first one is to clarify targets according to an existing classification, or according to new classifications. The second one is to re-express aircraft behavior in term of an existing tactical plan. The third one is to complete a sketchy current model of the pilot intents. if any pilot's new tactics,

then include them. The fourth one is to take into account individual differences between pilots: from a general model, tailor a specific one for a specific pilot.
She insists also on the necessity to be able to aggregate data, to display the information in way significant to the pilot, to be able to diagnose system malfunctions, and to coordinate emergency procedures.

Michalski is quite optimistic and thinks that relatively little effort is necessary to put information his programs need to start learning rules. He insists on the fact that using techniques that allow the use of combination of existing descriptors (those given by the human expert) may lead to drastic simplification of the rules.

Mitchell declares that he could not figure out precisely enough what could be a pilot-assistant theory. He therefore thought rather about a (car) driver assistant.
Let organize the knowledge under the form of schemas and schemas taxonomies.
Example of a schema: *slow-down-early-to-get-through-quickly.*
Description of the schema. Slow down early, far behind other vehicles. When other vehicles start moving again, hold velocity constant. When other vehicles reach one's speed, accelerate.
Schemas to apply afterwards: *Normal-traffic*, etc ...
Explanation of schema: speeding from complete stop requires more time than coasting through.
Possible confusions to be avoided: *slow-down-to-avoid-pedestrian*, *slow-down-to-read-sign*, etc ...
Taxonomies of schemas:

```
        Negotiate-intersections                          Pass-car
          /        !        `                              /    !
         /         !         `                            /     !
    Stop-signs    ...     Traffic-lights                 ...   ...
```

Use of these schemas: give advice to pilot.
How to learn these schemas?: LEAP (Tom Mitchell), ARMS (Jerry DeJong), ESA (Jerry DeJong). Interpret complex set of data, parse the situation.
Role of the explanation: allows to take into account the interesting features of the situation.

Milestones

3 yr: prototype of the schema acquisition.
7 yr: prototype to flight simulator.
10 yr: first demo inboard.

Pat Langley's 10 (N) years plan.
2 yr: generate artificial data from which rules can be deduced.
5 yr: protocol analysis.
N yr: robust plan understanding.
10 yr: develop plan understander at the level of the state-of-the-art.

Jerry DeJong disagrees with Pat's pessimism about the level of the state-of-the-art in plan understanding.

John Laird 's view about emergency procedures.
Given:
- procedures for emergency, - model of effect of actions, - access to success & failure information, - data base of scenario
Obtain:
a way to handle multiple emergencies.

This is a learning-by-problem-solving problem: when a particular solution has been found, induce a general procedure.
Issues addressed.
Availability of domain theory, availability of detailed scenarios, solving problems of goal interaction and of temporal dependencies, learning of general symbolic representations, optimization of slow systems.